

Automated Guidelines Implemented via the World Wide Web

James J. Cimino, M.D., Socrates A. Socratous, and Paul D. Clayton, Ph.D.
Department of Medical Informatics
Columbia University
New York, New York

The Arden Syntax provides a standard language for representing Medical Logic Modules (MLMs). Clinical decision support systems make use of MLMs by executing them when a "triggering" situation is detected, such as the storage of a relevant laboratory result. MLMs, in turn, usually require some additional patient-specific data which must be obtained by querying the clinical information system. Depending on the outcome of the logic process, a message may be generated which is routed to the appropriate person (e.g., a care giver, a hospital administrator, or a system manager). MLMs may also be chained together to perform the complex logic found in clinical guidelines. In order for a clinical system to make use of MLMs for clinical guidelines, it must be able to: (a) trigger MLMs appropriately, (b) respond to information requests from MLMs, and (c) chain MLMs together.

Guidelines typically start with some basic patient information and use branching logic to request additional information as needed. Sometimes the message produced is simply a recommendation to obtain additional information (e.g., perform a test). In some cases, a care giver may have knowledge of the additional information but it may not be available to the clinical information system. In situations such as these, it is desirable for the care giver to interact directly with the guideline logic to "fill in the blanks." Adapting a clinical information system to support automated guidelines can be a complex task. We have developed a simple approach that takes advantage of the capabilities of the World Wide Web to implement guidelines which can be triggered as needed and can obtain information from the clinical system and the care giver.

We encode all of the logic of a guideline in a single program called a Common Gateway Interface (CGI) which can be called via a Web browser. The CGI takes parameters, processes the MLM logic and queries the clinical information system as needed. The CGI returns a document to the Web browser that is in Hypertext Markup Language (HTML) format. Based on the information available to the CGI and the result of the logic process, one of three types of documents may be returned.

First, if the logic has sufficient information, the guideline may be traversed to completion, generating a recommendation message. Second, the logic may only be able to reach some intermediate state, and a recommendation may be made to obtain further information about the patient. Third, the logic may require information which can be obtained through interaction with the care giver. This last type of message takes advantage of capabilities of the Web.

The document requesting additional information is actually a form, with which the care giver can answer one or more questions. Included in the HTML document is a reference to the original CGI. When the care giver enters the additional information, the CGI is called again. The parameter values which were originally provided to the CGI are included in the HTML document and are again provided to the CGI. This time, however, the information provided by the care giver is included in the parameter list. The CGI executes its logic again and returns a document based on the additional information. Once again, this document may request further information and the cycle repeats until a terminal node in the guideline logic is reached.

Our CGIs make use of a vocabulary server which is accessible through Remote Procedure Calls (RPCs) on the Internet. These RPCs can be used in a proactive or reactive manner. In the proactive mode, the vocabulary server is queried for all of the possible answers to the question being posed to the care giver. For example, if the logic needs to know if the patient has any risk factors for cardiovascular disease, the CGI can request all members of the class "Cardiovascular Risk Factors" and present the results in the HTML form for the care giver to select. In the reactive mode, a question might be answered with free-text entry and the CGI passes the text to the vocabulary server to perform a match.

The poster illustrates how the various components in the process interact and shows how a single, self-contained CGI, together with a vocabulary server, can function as an institution-independent, automated decision support resource on the Internet.